

Variable Selection in Bayesian Semiparametric Regression Models

Ofir Harari^a, David M. Steinberg^a

^a*Department of Statistics and Operations Research, Raymond and Beverly Sackler Faculty of Exact Sciences,
Tel Aviv University, Tel Aviv 69978, Israel*

Abstract

In this paper we extend existing Bayesian methods for variable selection in Gaussian process regression, to select both the regression terms and the active covariates in the spatial correlation structure. We then use the estimated posterior probabilities to choose between relatively few modes through cross-validation, and consequently improve prediction.

1. Introduction

Gaussian processes have been employed extensively both in regression and classification problems in machine learning (see Rasmussen and Williams 2006). Their earliest use in statistical modeling may have been for fitting Kriging predictors in geostatistics (see e.g. Krige 1951, Cressie 1993) and they have been the most popular metamodeling approach to computer simulations (see Santner et al. 2003, chapter 3).

The problem of variable selection in Gaussian process regression has been addressed by Linkletter et al. (2006), while Savitsky et al. (2011) extended their results to generalized linear nonparametric models. In both cases no regression terms were included, and dependence on the factors was modeled by the authors solely via the spatial correlation structure.

In many applications only interpolation or smoothing are required, and the omission of regression terms can simplify the fitting of the Kriging metamodel, with little expense in predictive precision (see e.g. Sacks et al. 1989). It is also known that these models have a strong “Kriging towards the mean” tendency at sites far away from the training data, and it is thus

Email addresses: ofirhara@post.tau.ac.il (Ofir Harari), dms@post.tau.ac.il (David M. Steinberg)

often essential to add a trend for effective extrapolation (see Journel and Rossi 1989 for the original argument). Currently, the common practice is to either include all of the covariates both in the correlation structure and the linear component (i.e. *Universal Kriging*) or to omit the linear part altogether and model the data as observations from a pure (weakly) stationary random process (i.e. *Ordinary Kriging*). Joseph et al. (2008) proposed a Bayesian procedure for designed computer experiments, to select first and second order polynomials for the mean surface while leaving the covariance terms untouched. In this work we allow for exclusion of covariates from the correlation structure, along with their possible inclusion in the trend component, in the case that their only contribution to the output is linear.

2. Model Assumptions

Let $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathcal{X}$ be our training set, where $\mathcal{X} \subset \mathbb{R}^p$ is the experimental region, and let $\mathbf{y} = [y(\mathbf{x}_1), \dots, y(\mathbf{x}_n)]^\top$ be a vector of observations from the model

$$y(\mathbf{x}) = \beta_0 + \mathbf{x}^\top \boldsymbol{\beta} + Z(\mathbf{x}) + \varepsilon(\mathbf{x}) \quad , \quad (1)$$

where β_0 is a constant, $\boldsymbol{\beta} = [\beta_1, \dots, \beta_p]^\top$ is a vector of regression coefficients, $Z(\mathbf{x})$ is a zero mean stationary Gaussian process with marginal variance σ_Z^2 and a separable correlation function of the form

$$R(\mathbf{u}, \mathbf{v}; \boldsymbol{\rho}) = \prod_{j=1}^p \rho_j^{(u_j - v_j)^2} \quad (2)$$

for some vector of correlation parameters $\boldsymbol{\rho} = [\rho_1, \dots, \rho_p]^\top$, $0 \leq \rho_j \leq 1$, $1 \leq j \leq p$, and $\varepsilon(\mathbf{x})$ is a white noise process, independent of $Z(\mathbf{x})$, with variance parameter σ_ε^2 . Note that the closer ρ_j is to 1, the less $Z(\mathbf{x})$ tends to vary in the j th direction, and vice versa. This point is illustrated in Figure 1, in which a single realization of $Z(\mathbf{x})$ is drawn, with $p = 2$, a relatively large value of ρ_1 and an extremely small value of ρ_2 .

Denoting $\lambda = \sigma_\varepsilon^2 / \sigma_Z^2$, it can be easily verified that

$$\mathbf{y} | \beta_0, \boldsymbol{\beta}, \sigma_Z^2, \boldsymbol{\rho}, \lambda \sim \mathcal{N} \{ \beta_0 \mathbf{1}_n + \mathbf{X} \boldsymbol{\beta}, \sigma_Z^2 [\mathbf{R}(\boldsymbol{\rho}) + \lambda \mathbf{I}] \} \quad , \quad (3)$$

where \mathbf{X} is the design matrix and

$$\mathbf{R}_{lm}(\boldsymbol{\rho}) = R(\mathbf{x}_l, \mathbf{x}_m) = \prod_{j=1}^p \rho_j^{(\mathbf{x}_{lj} - \mathbf{x}_{mj})^2} \quad .$$




Figure 1: A realization of $Z(\mathbf{x})$ with the correlation function (2), for $p = 2, \rho_1 = 0.75$ and $\rho_2 = 5 \times 10^{-6}$.

We can interpret (1) as follows: each of the p covariates in \mathbf{x} may (or may not) make a linear contribution to the response $y(\mathbf{x})$. If, in addition, it makes a nonlinear contribution, it should be captured by $Z(\mathbf{x})$, whose flexible nature makes it free of any rigid structure. The white noise term, also known as the “nugget” (see Andrianakis and Challenor 2012) and interpreted as “measurement error”, is often omitted when using the Kriging metamodel to emulate deterministic computer simulations. However, its (possibly unnatural) inclusion has certain benefits, both numerical and in terms of the average squared prediction error. The consequent λ term in (3) is mostly identified with the equivalent penalized least squares minimization problem in the Reproducing Kernel Hilbert Space corresponding to $R(\cdot, \cdot)$ (see Wahba 1990). In this paper, we would like to examine the possibility of leaving some of the covariates out of the correlation function. By doing that, we deliberately omit variables that are (in the case of computer experiments) known for a fact to be part of the computer code. It is then reasonable to add an error term, to avoid interpolation at the learning set, in favor of better predictions at new sites.

3. Formalizing as a Bayesian Model

We would now like to turn our attention to the problem of variable selection in model (1). Research on Bayesian variable selection in linear regression models has been thoroughly

documented - Chipman et al. (2001) give a detailed account of some theoretical and computational aspects Linkletter et al. (2006) and Savitsky et al. (2011) apply point-mass mixture prior distributions for the correlation parameters of simple Kriging models. In this work, we aim to propose a unified framework, in which “spike-and-slab” priors, like those used by Linkletter et al. (2006) and Savitsky et al. (2011), are employed to select the components of both the regression and spatial parts of the model.

3.1. Writing the Posterior Distribution

To put the above to practice, we now assume that the regression coefficients are independent, with prior distributions

$$\beta_j | \omega_r \sim \omega_r \mathcal{N}(0, \tau^2) + (1 - \omega_r) \delta_0, \quad (4)$$

where δ_0 is a point mass at zero. Alternatively, one can denote by $\gamma^r = (\gamma_1^r, \dots, \gamma_p^r)$ a vector of latent variables, in which $\gamma_j^r = \mathbb{1}\{\beta_j \neq 0\}$ indicates whether or not the j th covariate is included in the linear component. We may then write $\gamma_j^r | \omega_r \sim \text{Binom}(1, \omega_r)$ and $\beta_j | \gamma_j^r \sim \mathcal{N}(0, \gamma_j^r \tau^2)$.

Regarding the spatial component, for the j th variable to be inert, its corresponding correlation factor ρ_j must be equal to 1. An analogue of (4) would then be

$$\rho_j | \omega_c \sim \omega_c \mathcal{U}(0, 1) + (1 - \omega_c) \delta_1, \quad (5)$$

implemented by defining the vector of latent variables $\gamma^c = (\gamma_1^c, \dots, \gamma_p^c)$, where $\gamma_j^c | \omega_c \sim \text{Binom}(1, \omega_c)$ and

$$\pi(\rho_j | \gamma_j^c) = \gamma_j^c \mathcal{U}(0, 1) + (1 - \gamma_j^c) \delta_1(\rho_j).$$

Using Bayes rule, we may now write the posterior distribution

$$\begin{aligned} & p(\gamma^r, \gamma^c, \omega_r, \omega_c, \beta_0, \beta, \rho, \sigma_z^2, \lambda | \mathbf{y}) \\ & \propto p(\mathbf{y} | \beta_0, \beta, \sigma_z^2, \rho, \lambda) \pi(\gamma^r, \beta | \omega_r) \pi(\gamma^c, \rho | \omega_c) \times \pi(\beta_0, \sigma_z^2, \lambda, \omega_r, \omega_c), \end{aligned} \quad (6)$$

where $p(\mathbf{y} | \beta_0, \beta, \sigma_z^2, \rho, \lambda)$ is the multivariate Normal density, as indicated in (3),

$$\pi(\gamma^r, \beta | \omega_r) = \omega_r^{\sum_j \mathbb{1}\{\gamma_j^r=1\}} (1 - \omega_r)^{\sum_j \mathbb{1}\{\gamma_j^r=0\}} \prod_{\gamma_j^r=1} \phi_\tau(\beta_j)$$

(here $\phi_\tau(\beta_j)$ is the Normal pdf with zero mean and standard deviation τ , evaluated at β_j) and

$$\pi(\gamma^c, \rho | \omega_c) = \omega_c^{\sum_j \mathbb{1}\{\gamma_j^c=1\}} (1 - \omega_c)^{\sum_j \mathbb{1}\{\gamma_j^c=0\}}.$$

3.2. Markov Chain Monte Carlo Sampling from the Posterior Distribution

We now apply a Metropolis Sampling algorithm to draw a random sample from the posterior distribution. A suitable ‘jump’ (or ‘proposal’) distribution is required - one that will allow the resulting Markov chain to arrive at any possible state with positive probability. Here we extend the key ideas of Savitsky et al. (2011), to propose a full Metropolis-Hastings schema:

1. Update (γ^r, β) and (γ^c, ρ) : draw $k \in \{0, 1, \dots, 2p\}$ from a $\text{Binom}(2p, \nu)$ distribution (typically $\nu = (2p)^{-1}$) and change the present state of k randomly chosen latent variables out of (γ^r, γ^c) .
 - If the value of γ_j^r changes from 0 to 1, sample $\tilde{\beta}_j$ from a $\mathcal{N}(0, \tau^2)$ proposal.
 - If the value of γ_j^c changes from 0 to 1, sample $\tilde{\rho}_j$ from a $\mathcal{U}(0, 1)$ proposal.
 - If the value of γ_j^r changes from 1 to 0, set $\tilde{\beta}_j = 0$. This results in covariate \mathbf{x}_j being omitted from the linear trend in the current iteration.
 - If the value of γ_j^c changes from 1 to 0, set $\tilde{\rho}_j = 1$. This results in covariate \mathbf{x}_j being omitted from the correlation function in the current iteration.
- * In any of the above cases, add a small jitter to the remaining β and ρ values (while making sure all the ρ_j -s remain in the interval $(0, 1)$).
2. Update $(\beta_0, \sigma_z^2, \lambda, \omega_r, \omega_c)$: to simplify matters, transform the parameters to

$$\mu = \log(\sigma_z^2) \quad ; \quad \zeta = \log(\lambda) \quad ; \quad \psi_r = \log \frac{\omega_r}{1 - \omega_r}$$

and

$$\psi_c = \log \frac{\omega_c}{1 - \omega_c} \quad ,$$

and sample $(\tilde{\beta}_0, \tilde{\mu}, \tilde{\zeta}, \tilde{\psi}_r, \tilde{\psi}_c)$ from a multivariate Normal distribution about the present state, $(\beta_0, \mu, \zeta, \psi_r, \psi_c)$. In our implementation we used independent components. Remember to take the Jacobian term

$$\frac{\partial(\sigma_z^2, \lambda, \omega_r, \omega_c)}{\partial(\mu, \zeta, \psi_r, \psi_c)} = \mu \zeta \omega_r (1 - \omega_r) \omega_c (1 - \omega_c)$$

into account in the posterior later.

3. The proposed vector $(\tilde{\gamma}^r, \tilde{\beta}, \tilde{\gamma}^c, \tilde{\rho}, \tilde{\beta}_0, \tilde{\sigma}_z^2, \tilde{\lambda}, \tilde{\omega}_r, \tilde{\omega}_c)$ is accepted with probability

$$\alpha = \min \left\{ 1, \frac{p(\tilde{\gamma}^r, \tilde{\gamma}^c, \tilde{\beta}_0, \tilde{\beta}, \tilde{\rho}, \tilde{\sigma}_z^2, \tilde{\lambda}, \tilde{\omega}_r, \tilde{\omega}_c | \mathbf{y})}{p(\gamma^r, \gamma^c, \beta_0, \beta, \rho, \sigma_z^2, \lambda, \omega_r, \omega_c | \mathbf{y})} \right\} . \quad (7)$$

Note that (7) holds, since by construction, the proposal distribution described above is symmetric, i.e.

$$p(\tilde{\gamma}^r, \tilde{\beta}, \tilde{\gamma}^c, \tilde{\rho} | \gamma^r, \beta, \gamma^c, \rho) = p(\gamma^r, \beta, \gamma^c, \rho | \tilde{\gamma}^r, \tilde{\beta}, \tilde{\gamma}^c, \tilde{\rho})$$

and

$$p(\tilde{\beta}_0, \tilde{\sigma}_z^2, \tilde{\lambda} | \beta_0, \sigma_z^2, \lambda) = p(\beta_0, \sigma_z^2, \lambda, \omega_r, \omega_c | \tilde{\beta}_0, \tilde{\sigma}_z^2, \tilde{\lambda}, \tilde{\omega}_r, \tilde{\omega}_c) .$$

4. Prediction Strategies

We now wish to take advantage of the procedure described in Section 3.2, to make accurate predictions at m new sites \mathbf{X}_{new} . We will propose two possible approaches, the first of which - based on model averaging - is a more “orthodox” Bayesian approach, while the second one, in which we merely use the schema of Section 3.2 to identify the MAP model, might perhaps be characterized as opportunistic.

4.1. Prediction based on Model Averaging

Conditional on $\Theta = (\beta_0, \beta, \rho, \sigma_z^2, \lambda)$, the joint distribution of $\mathbf{y}_{\text{old}} = y(\mathbf{X}_{\text{old}})$ and $\mathbf{y}_{\text{new}} = y(\mathbf{X}_{\text{new}})$ is given by

$$\begin{bmatrix} \mathbf{y}_{\text{new}} \\ \mathbf{y}_{\text{old}} \end{bmatrix} = \mathcal{N} \left\{ \begin{bmatrix} \beta_0 \mathbf{1}_m + \mathbf{X}_{\text{new}} \beta \\ \beta_0 \mathbf{1}_n + \mathbf{X}_{\text{old}} \beta \end{bmatrix} ; \Sigma \right\} ,$$

where

$$\Sigma = \sigma_z^2 \left[\begin{array}{c|c} \mathbf{R}_{\text{new}} + \lambda \mathbf{I}_m & \mathbf{R}_{\text{new,old}} \\ \hline \mathbf{R}_{\text{new,old}}^\top & \mathbf{R}_{\text{old}} + \lambda \mathbf{I}_n \end{array} \right]$$

and $\mathbf{R}_{\text{new,old}}$ is the cross-correlation matrix

$$\mathbf{R}_{\text{new,old}}^{i,j} = R(\mathbf{x}_{\text{new}}^i, \mathbf{x}_{\text{old}}^j) .$$

It is therefore straightforward to show that

$$\begin{aligned} \mathbb{E} \left\{ \mathbf{y}_{\text{new}} \middle| \mathbf{y}_{\text{old}}, \boldsymbol{\Theta} \right\} &= \beta_0 \mathbf{1}_m + \mathbf{X}_{\text{new}} \boldsymbol{\beta} \\ &+ \mathbf{R}_{\text{new,old}} [\mathbf{R}_{\text{old}} + \lambda \mathbf{I}_n]^{-1} (\mathbf{y}_{\text{old}} - \beta_0 \mathbf{1}_n - \mathbf{X}_{\text{old}} \boldsymbol{\beta}) \quad , \end{aligned} \quad (8)$$

and predictions, based on a sample $\{\boldsymbol{\Theta}_i\}_{i=1}^N$ from the posterior distribution (as described in section 3.2), can be made by integrating $\boldsymbol{\Theta}$ out, namely

$$\begin{aligned} \hat{\mathbf{y}}_{\text{new}} &= \mathbb{E} \left\{ \mathbf{y}_{\text{new}} \middle| \mathbf{y}_{\text{old}} \right\} = \int \mathbb{E} \left\{ \mathbf{y}_{\text{new}} \middle| \mathbf{y}_{\text{old}}, \boldsymbol{\Theta} \right\} p(\boldsymbol{\Theta} | \mathbf{y}_{\text{old}}) d\boldsymbol{\Theta} \\ &\approx \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left\{ \mathbf{y}_{\text{new}} \middle| \mathbf{y}_{\text{old}}, \boldsymbol{\Theta}_i \right\} \quad . \end{aligned} \quad (9)$$

This is, in fact, the usual model averaging estimator (see e.g. Ghosh et al. 2006, Chapter 9). If the set \mathbf{X}_{new} is known in advance, (8) can be calculated on the fly (within each Metropolis iteration), and significant calculation time can be saved, depending on the acceptance rate of our sampling - provided that no new predictions are made whenever a new vector is rejected. It is also worthwhile to consider “denoising” through elimination of highly unlikely models from (9), by setting exceptionally small posterior probabilities (according to a predetermined threshold) to zero.

4.2. Prediction based on Variable Selection

Another obvious option is to regard the Bayesian mechanism of Section 3 as a mere selection procedure to identify influential (either linear or nonlinear) variables. In that case, inert variables in either the regression or the correlation part, will be set to 0 or 1, respectively, and the prediction at location \mathbf{x} will be

$$\hat{y}(\mathbf{x}) = \mathbf{x}^\top \hat{\boldsymbol{\beta}} + \mathbf{r}_{\hat{\boldsymbol{\rho}}}^\top(\mathbf{x}) \left[\mathbf{R}_{\hat{\boldsymbol{\rho}}} + \hat{\lambda} \mathbf{I} \right]^{-1} (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}}) \quad , \quad (10)$$

where $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\rho}}, \hat{\lambda})$ are the maximum likelihood estimators of $(\boldsymbol{\beta}, \boldsymbol{\rho}, \lambda)$, obtained by solving iteratively

$$\begin{cases} \hat{\boldsymbol{\beta}}(\hat{\boldsymbol{\rho}}, \hat{\lambda}) = \left(\mathbf{X}^\top [\mathbf{R}_{\hat{\boldsymbol{\rho}}} + \hat{\lambda} \mathbf{I}]^{-1} \mathbf{X} \right)^{-1} [\mathbf{R}_{\hat{\boldsymbol{\rho}}} + \hat{\lambda} \mathbf{I}]^{-1} \mathbf{y} \quad , \\ \hat{\sigma}^2(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\rho}}, \hat{\lambda}) = n^{-1} (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}})^\top [\mathbf{R}_{\hat{\boldsymbol{\rho}}} + \hat{\lambda} \mathbf{I}]^{-1} (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}}) \quad , \\ (\hat{\boldsymbol{\rho}}, \hat{\lambda}) = \underset{\boldsymbol{\rho}, \lambda}{\operatorname{argmin}} \left\{ n \log \hat{\sigma}^2(\hat{\boldsymbol{\beta}}(\boldsymbol{\rho}, \lambda), \boldsymbol{\rho}, \lambda) + \log |\mathbf{R}_{\boldsymbol{\rho}} + \lambda \mathbf{I}| \right\} \quad . \end{cases}$$

Details of the optimization routine are given in Roustant et al. (2012).

Denoting $\mathcal{M} = (\gamma^r, \gamma^c)$, our model of choice could then be the highest posterior probability model, i.e.

$$\mathcal{M}^* = \underset{\mathcal{M}}{\operatorname{argmax}} p(\mathcal{M}|\mathbf{y})$$

(to which we will refer in the text as the *maximum a posteriori model*, or simply - the *MAP model*), or, alternatively, we may set a threshold q on the *posterior inclusion probability*

$$p_i = \sum_{\ell: \gamma_i=1} p(\mathcal{M}_\ell|\mathbf{y}) \quad ,$$

(where γ_i may be either γ_i^r or γ_i^c), and select the variable x_i to either part, only if $p_i \geq q$. The choice of $q = \frac{1}{2}$ leads to what is famously known as the ‘median model’, which - in the simpler case of linear regression with uncorrelated noise - has been shown by Barbieri and Berger (2004) to be the best predictive model, under very strict conditions on both the prior distributions and the covariates where predictions are to be made. In Section 5 we set the threshold on the posterior inclusion probability at a relatively high value of 0.8. A more educated choice of that value is proposed in Subsection 6.1, based on v-fold cross-validation.

5. A Simulation Study

We now wish to test the proposed methods on simulated data. For that purpose, we used a 35-run maximin Latin hypercube design (see Jones and Johnson 2009) in $[-0.75, 0.75]^5$ to sample noisy data from the allegedly 5 dimensional function

$$y(\mathbf{x}) = 3X_2 + 4X_3 + 5X_4 + 5 \cos \frac{3\pi X_1}{2} + 4 \cos \frac{2\pi X_2}{2} + 3 \cos \frac{\pi X_3}{2} + \varepsilon(\mathbf{x}) \quad , \quad (11)$$

where $\varepsilon(\mathbf{x}) \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 0.1^2)$. Note that while (11) consists of both linear and non-linear parts, the variables X_1 and X_4 are absent from the linear and non-linear parts, respectively, and X_5 is inert. Figure 2 shows how the different (additive) components of function (11) vary as functions of their respective variables.

We assigned prior distributions as follows:

$$\beta_j | \gamma_j^r = 1 \sim \mathcal{N}(0, 5^2) \quad , \quad j = 1, \dots, 5 \quad ; \quad \omega_r, \omega_c \sim \mathcal{U}(0, 1) \quad ; \quad \beta_0 \sim \mathcal{N}(0, 10^2) \quad ;$$

$$\sigma_z^2 \sim \text{Inv-}\Gamma(3, 2) \quad \text{and} \quad \lambda \sim \text{Inv-}\Gamma(3, 0.2) \quad ,$$




Figure 2: Plotting function (11) against each of $X_1 - X_4$.

and calculated the root mean square prediction error (RMSPE) on an independent 100-run maximin LHD in the same region. A size 250,000 sample was drawn from the posterior distribution, as described in Section 3.2, after the first 5000 were discarded as burn-in.




Figure 3: Box plots of the posterior distributions of the regression and correlation parameters, for the simulated data.

Figure 3 shows the posterior distributions of both the regression coefficients and the correlation parameters, while Figure 4 shows their respective marginal posterior inclusion probabilities. For now, we set the inclusion limit at 0.8, leaving us with a single regression term (X_4) in the trend part of the model, and three covariates in the nonlinear component (X_1 , X_2 and X_3). This is also consistent with Figure 3. Apparently, the Gaussian process term effectively reflects the linear effects of X_2 and X_3 as well as their nonlinear effects, which explains why they were dropped from the regression component. As expected, Figure 3 clearly identifies X_1 as inactive in the trend part and X_4 in the non-linear component, while X_5 is deemed inert beyond any doubt.




Figure 4: The posterior inclusion probabilities for the simulated data. The dashed line marks 80% posterior inclusion frequency.




Figure 5: Predictions vs. true values for the simulated example.

Table 1 contains a comparison of the empirical RMSPE of several fitted models when tested on the 100-run validation set. The models include ordinary kriging (OK, no linear trend at all), universal kriging (UK, a linear trend which includes all of the covariates), model averaging based on (9), variable selection based on the posterior inclusion probability (0.8 at least) and the empirical highest posterior probability (MAP) model.

Table 1: Comparison of several prediction methodologies for the simulated data.

| Method | RMSPE |
|---------------------|-------|
| OK | 0.168 |
| UK | 0.130 |
| Averaging | 0.247 |
| Posterior Inclusion | 0.115 |
| MAP | 0.124 |

Interestingly, the MAP model in this example was the one containing X_2 , X_3 and X_4 in its linear part, along with X_1 , X_2 and X_3 in its non-linear part, compatible with the mechanism generating the data. However, it was outperformed by the model selected according to the posterior inclusion probabilities. It is our recommendation to resist the temptation to blindly opt for the empirical MAP model, when it is unclear whether or not it is the true highest posterior model - in particular when the posterior distribution is multimodal or flat near its mode. In this example, not for the only time in this paper, model averaging proved to be a poor strategy. One possible reason could be poor estimation of the posterior distribution, resulting in excessive weight on low priority models. Figure 5 demonstrates the performance on the hold-out set of the model selected by the posterior inclusion probabilities.

6. Applications

6.1. Heat Exchanger Emulator

We now wish to test our method on the computer simulation described in Qian et al. (2006). Here, the total rate of steady state heat transfer of a heat exchanger for electronic cooling applications is modelled as a function of four input factors: the mass flow rate of entry air, the temperature of entry air, the solid material thermal conductivity and the temperature of the heat source. The response was evaluated at a 64-run experimental design, constituting an orthogonal array-based LHD, and produced by a finite difference code. An additional 14 run test set was used to validate the resulting emulator. Input variables were standardized into the $[0, 1]$ region before analysis, to overcome the different scales. Using the prior distributions of Section 5, a sample of size 155,000 was drawn from the posterior distribution, with the first 5000 serving as burn-in.

Previously, only interpolators have been attempted, as the simulation itself is deterministic. Qian et al. (2006) reported a RMSPE of 2.59 using a universal Kriging model with a linear trend, compared to 5.15 achieved by ordinary Kriging, while Ba and Joseph (2012) reduced the RMSPE to 2.24. Lately, Harari and Steinberg (2014a) managed an RMSPE of 1.93, using a convex combination of two independent Gaussian processes to model the spatial structure. Here we allowed for a non-zero nugget effect. The different methods are compared in Table 2. One interesting feature is that x_4 has a strong linear effect but is slightly less dominant in the covariance. The other factors are modeled primarily via the covariance. This slightly

deviates from the work of Qian et al. (2006), who used linear terms for x_2 and x_4 , and modelled the correlation through all four factors.

The dashed, horizontal line in Figure 7 marks the cutoff used in this example to omit terms from the model. The cutoff - 77% - was obtained via the following procedure: only terms with at least a 30% marginal inclusion probability were considered, and terms with 90% marginal inclusion probability or more were automatically included. That left 4 competing models, of which the eventual model was chosen via 8-fold cross-validation (fitting the model, based on 56 observations and testing on the remaining 8 each time). Figure 11 shows the results, along with one-standard error marks, potentially serving to prioritize sparse models, in a similar manner to CART models (see Breiman et al. 1984, chapter 3). Generally, using this procedure leaves us with at most 2^p models to compare (and, in practice, typically far less), instead of 2^{2p} in an exhaustive search. As evident from Figure 8, in this example, a model consisting of 4 predictors (X_1 in the linear part and X_2 , X_3 and X_4 in the stochastic part), was selected.




Figure 6: Box plots of the posterior distributions of the regression and correlation parameters for the heat exchanger data.

6.2. The Boston Housing Data

To test our method on a larger-scale problem, we now analyze the Boston Housing data set, which also appeared in Savitsky et al. (2011), and - most famously - in Breiman and Friedman (1985), where a full description of the 13 predictors - related to the median value of owner-occupied homes in census tracts of the Boston metropolitan area - can be found. The data




Figure 7: The posterior inclusion probabilities for the heat exchanger simulation. The dashed line marks a 77% posterior inclusion level.




Figure 8: 8-fold cross-validation results for the Heat Exchanger simulation data, including 1-SE marks.

Table 2: Comparison of several prediction methodologies for the heat exchanger example.

| Method | RMSPE |
|---------------------|-------|
| OK | 2.023 |
| UK | 2.670 |
| Averaging | 2.686 |
| Posterior Inclusion | 1.793 |
| MAP | 1.793 |

set contains 506 observations, of which we used 130 as a training set and the other 376 for validation. Again, we transformed the predictors so that they are all in the $[0, 1]$ interval. We produced a sample from the posterior distribution of size 250,000 (with an additional burn-in sample of size 15,000), and the same set of priors as in Section 5. The resulting marginal

posterior distributions are shown in Figure 9, where it is evident that none of the covariates makes a pure, linear contribution to the median property value, and that the 6th (average number of rooms per dwelling), 10th (full-value property-tax rate per \$10,000), 11th (pupil-teacher ratio by town) and 13th (% lower status of the population) covariates stand out in the non-linear part of the model.




Figure 9: Box plots of the posterior distributions of the regression and correlation parameters, for the Boston Housing data.




Figure 10: The posterior inclusion probabilities for the Boston Housing data set. The dashed line marks a 76% posterior inclusion frequency.

Figure 10, where the marginal posterior inclusion probabilities are shown, is very much in line with Figure 9. As previously in 6.1, the cutoff - 66% - was determined through a 10-fold cross-validation procedure, shown in Figure 11. In this example, the most parsimonious

model, consisting of only 4 predictors, was singled out as the obvious choice.




Figure 11: 10-fold cross-validation results for the Boston Housing data, including 1-SE marks.

In Figure 12, the reader can see how the emulator values $\hat{y}(\mathbf{x})$ vary for each of the covariates. We evaluated (10) at numerous sites on a regular grid, and the plotted values in each graph were obtained by averaging on the other three variables. It should come as no surprise to the reader, based on Figure 9, that X_6 and X_{13} are the variables accounting for the dominant nonlinear effects, while the output seems to be far less sensitive to changes in the values of X_{10} and X_{11} , consequence of their $\hat{\rho}$ estimates of 0.392 and 0.668, respectively. Thorough discussion of sensitivity analysis can be found in Saltelli et al. (2008). While the reader may be misled by Figure 12 to think that X_{10} and X_{11} are linear effects in this example, we feel the need to emphasize that these plots were obtained by averaging, and do not capture interactions. That also explains the flat (or even slightly decreasing) average behavior of $\hat{y}(\mathbf{x})$ for the lower values of X_6 (average number of rooms per dwelling).

It is interesting to note that while the ordinary kriging and universal kriging models achieved empirical RMSPEs of 4.02 and 4.01, respectively, on the hold-out set, the selected model, based on 4 predictors, achieved a RMSPE of 3.98. Although prediction is not always the main purpose of variable selection procedures of this sort, it is still encouraging to find out that dimensionality reduction of this magnitude does not necessarily come at the expense of predictive accuracy. For further reference, we used the Treed Gaussian Process (TGP) model, a nonstationary model, developed and implemented by Gramacy and Lee (2008), which combines partitioning of the predictor space with local, stationary GP models. In this example, the TGP model achieved a RMSPE of 4.16.




Figure 12: The predicted value $\hat{y}(\mathbf{x})$ for the Boston Housing data, plotted against X_6 , X_{10} , X_{11} and X_{13} , averaged over all other variables in each plot.

7. Discussion

In this work, we explored variable selection in Bayesian semiparametric regression models. Although the primary goal of these procedures is often of qualitative or explanatory nature, we keep one eye open on the accuracy of our predictions. With that in mind, our variable selection procedure is driven by prediction throughout, as evident from the cross-validation procedure described in Subsections 6.1 and 6.2, and unlike the frequentist-like selection procedure of Linkletter et al. (2006), where identification of the active inputs was the sole purpose.

Ideally, we would like to take advantage of the Bayesian framework to provide accurate predictions, using (9). However, using model averaging to a good effect might require an educated choice of prior distributions, along with possibly vast samples from the posterior distributions. One very interesting open question is whether a set of sparsity prior distributions exists, such that under some aggregation methods, the minimax convergence rate of the mean squared prediction error can be achieved, as in the exponential weighting scheme of Rigollet and Tsybakov (2012) for linear regression models.

Naturally, one may consider replacing model (1) with universal kriging models of the form $y(\mathbf{x}) = \beta_0 + \mathbf{f}(\mathbf{x})^\top \boldsymbol{\beta} + Z(\mathbf{x}) + \varepsilon(\mathbf{x})$, where $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), \dots, f_r(\mathbf{x})]^\top$ for some dictionary of regression functions. In that case, extra caution needs to be taken, as choice of regression functions already in the reproducing kernel Hilbert space associated with the covariance kernel $R(\cdot, \cdot)$ could cause a multicollinearity-like effect (see Bursztyn and Steinberg 2004), resulting in the omission of crucial factors from our model.

One question left unanswered in this paper, and which may be the basis for future research, is whether or not one should include some linear terms for extrapolation purposes, even when the selection process indicates otherwise. Based on the artificial example of Section 5, it looks very likely that the model will drop most of the regression functions and rely on the covariance kernel. It would therefore be interesting to see how devastating an effect that would have on predictions outside the convex hull of the training data. Another possibility is to begin by “orthogonalizing” the covariance kernel to linear terms, by subtracting out the initial components of its Mercer expansion. Using Mercer’s expansion in computer experiments has been covered by Harari and Steinberg (2014b) and Fedorov and Flanagan (1998), among others.

Supplementary Material

R code and data files: R code for all sections and data files can be downloaded by the readers as a .zip file.

References

- Andrianakis, I., Challenor, P., 2012. The Effect of the Nugget on Gaussian Process Emulators of Computer Models. *Computational Statistics and Data Analysis* 56 (12), 4215–4228.
- Ba, S., Joseph, V., 2012. Composite Gaussian Process Models for Emulating Expensive Functions. *Annals of Applied Statistics* 6 (4), 1838–1860.
- Barbieri, M., Berger, J., 2004. Optimal Predictive Model Selection. *The Annals of Statistics* 32 (3), 870–897.
- Breiman, L., Friedman, J., 1985. Estimating Optimal Transformations for Multiple Regression and Correlation. *Journal of the American Statistical Association* 80 (391), 580–598.
- Breiman, L., Friedman, J., Olshen, R., Stone, C., 1984. *Classification and Regression Trees*. Statistics/Probability Series. Wadsworth Publishing Company, Belmont, California, U.S.A.
- Bursztyn, D., Steinberg, D., 2004. Data Analytic Tools for Understanding Random Field Regression Models. *Technometrics* 46 (4), 411–420.

- Chipman, H., George, E., McCulloch, R., 2001. The Practical Implementation of Bayesian Model Selection. Institute of Mathematical Statistics, Beachwood, OH, pp. 65–116.
- Cressie, N., 1993. Statistics for Spatial Data. Wiley-Interscience.
- Fedorov, V., Flanagan, D., 1998. Optimal Monitoring Network Design Based on Mercer's Expansion of Covariance Kernel. *Journal of Combinatorics, Information & System Sciences* 23 (1-4), 237–250.
- Ghosh, J., Delampady, M., Samanta, T., 2006. An Introduction to Bayesian Analysis: Theory and Methods. Springer.
- Gramacy, R., Lee, H., 2008. Bayesian Treed Gaussian Process Models With an Application to Computer Modeling. *Journal of the American Statistical Association* 103 (483), 1119–1130.
- Harari, O., Steinberg, D., 2014a. Convex Combination of Gaussian Processes for Bayesian Analysis of Deterministic Computer Experiments. *Technometrics* (to appear).
- Harari, O., Steinberg, D., 2014b. Optimal Designs for Gaussian Process Models via Spectral Decomposition. *Journal of Statistical Planning and Inference* (to appear).
- Jones, B., Johnson, R., 2009. The Design and Analysis of the Gaussian Process Model. *Quality and Reliability Engineering International* 25 (5), 515–524.
- Joseph, V., Hung, Y., Sudjianto, A., 2008. Blind Kriging: A New Method for Developing Metamodels. *Journal of Mechanical Design* 130 (3), 031102.
- Journel, A., Rossi, M., 1989. When Do We Need a Trend Model in Kriging? *Mathematical Geology* 21 (7), 715–739.
- Krige, D., 1951. A statistical Approach to some Mine Valuation and Allied Problems on the Witwatersrand. Master's thesis, University of the Witwatersrand.
- Linkletter, C., Bingham, D., Hengartner, N., Higdon, D., Ye, K., 2006. Variable Selection for Gaussian Process Models in Computer Experiments. *Technometrics* 48 (4), 478–490.
- Qian, P., Seepersad, C., Joseph, V., Allen, J., Wu, C., 2006. Building Surrogate Models with Detailed and Approximate Simulations. *ASME Journal of Mechanical Design* 128 (4), 668–677.

- Rasmussen, C., Williams, C., 2006. Gaussian Processes for Machine Learning. The MIT Press.
- Rigollet, P., Tsybakov, A., 2012. Sparse Estimation by Exponential Weighting. *Statistical Science* 27 (4), 558–575.
- Roustant, O., Ginsbourger, D., Deville, Y., 2012. DiceKriging, DiceOptim: Two R Packages for the Analysis of Computer Experiments by Kriging-Based Metamodeling and Optimization. *Journal of Statistical Software* 51 (1), 1–55.
- Sacks, J., Schiller, S., Welch, W., 1989. Designs for Computer Experiments. *Technometrics* 31 (1), 41–47.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Trantola, S., 2008. *Global Sensitivity Analysis: the Primer*. John Wiley.
- Santner, T., Williams, B., Notz, W., 2003. *The Design and Analysis of Computer Experiments*. Springer.
- Savitsky, T., Vannucci, M., Sha, N., 2011. Variable Selection for Nonparametric Gaussian Process Priors: Models and Computational Strategies. *Statistical Science* 26 (1), 130–149.
- Wahba, G., 1990. *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics.